



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/646,851	08/22/2003	Rui Liang	10830.0097NP	1500
27927	7590	08/04/2008	EXAMINER	
RICHARD AUCITERLONIE			KEEFER, MICHAEL E	
NOVAK DRUCE & QUIGG, LLP			ART UNIT	PAPER NUMBER
1000 LOUISIANA				2154
53RD FLOOR				
HOUSTON, TX 77002				
MAIL DATE		DELIVERY MODE		
08/04/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/646,851

Filing Date: August 22, 2003

Appellant(s): LIANG ET AL.

Richard C. Auchterlonie
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 5/23/2008 appealing from the Office action
mailed 12/10/2007.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

WITHDRAWN REJECTIONS

The following grounds of rejection are not presented for review on appeal because they have been withdrawn by the examiner. The rejection of claims 7-11 and 19-23 under 35 U.S.C. 112 second paragraph..

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

6564229	Baweja et al.	5-2003
2002/0010665	Lefebvre et al.	01-2002
5819292	Hitz et al.	10-1998
20020095616	Busser, Richard W	07-2002
6934804	Hashemi	08-2005
7010553	Chen et al.	3-2006
20010052021	Bolosky et al.	12-2001
7107385	Rajan et al.	09-2006

IANA. "Port Numbers"

Retrieved from

"<http://www.iana.org/assign>

ments/port-numbers" on

7/30/2008. Last updated on

7/28/2008.

PCMAG.COM. "Definition

of Unix Socket". Retrieved

from

"http://www.pcmag.com/encyclopedia_term/0,2542,t= Unix+socket&i=53469,00.asp

" on 7/30/2008.

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 27, 30, and 34-36 are rejected under 35 U.S.C. 102(e) as being anticipated by Chen et al. (US 7076509), hereafter Chen.

Regarding **claim 27**, Chen (US 7076509) discloses:

A network file server (Fig. 6, 600) comprising:

data storage; (Fig. 6, Disks connected to Disk driver 650)

an interface for coupling the data storage to a data network; (Fig. 6,

Media access 610) and

at least one processor programmed for permitting clients in the data network to access the data storage in accordance with a plurality of access protocols; (inherent, as multiple protocols are disclosed (Col. 9

lines 12-18, therefore a processor must be present and programmed to allow clients to use those protocols)

the data storage containing at least one file for storing file attributes and for storing metadata defining a virtual direct access storage device and for storing data of the virtual direct access storage device; (Fig. 6, File system 665 and/or VDISK module 670)

the access protocols including at least one block level access protocol for access to the virtual direct access storage device by accessing the metadata and data of the virtual direct access storage device (Fig. 6, FC 630, iSCSI 628); and

the access protocols including at least one file access protocol for accessing said at least one file. (Fig. 6, NFS 620, CIFS 622)

wherein the metadata includes attributes of the virtual direct access storage device, and the attributes of the virtual direct access storage device and the data of the virtual direct access storage device are stored together in a single file. (Col. 10, lines 43-47 states that a vdisk is a special file (singular) type.)

wherein the metadata includes attributes of the virtual direct access storage device, and the attributes of the virtual direct access storage device and the data of the virtual direct access storage device are stored in separate files in a common file system. (Col 10 lines 47-56 describe that

an embodiment of a vdisk may have one file (stream) containing metadata and other files in the file system (streams) would have data.)

wherein the attributes define an internal organization of the virtual direct access storage device and those attributes are stored in the single file. (Fig. 5 of fully incorporated application 10/216453 (now US 7107385) shows that the metadata does give an internal organization of the VLUN, showing where particular data is actually stored on a physical device.)

Regarding **claim 30 as applied to claim 24**, Chen (US 7076509) discloses:

which includes a snapshot copy facility for copying the data of the virtual direct access storage device concurrent with a client using the block level access protocol over the network to write data to the virtual direct access storage device. (Col. 13 lines 43-46 disclose a snapshot system)

Regarding **claim 34 as applied to claim 24**, Chen (US 7076509) discloses: wherein the network is an IP network, and the block level access protocol is SCSI. (Col 9, lines 56-59, note that iSCSI, an implementation of a block level access protocol (SCSI) over TCP (TCP/IP, or an IP network)

Regarding **claims 35-36 as applied to claim 24**, Chen (US 7076509) discloses: wherein the file access protocol is NFS. (Col. 9 lines 15-18, state that a client using the UNIX (or Linux) operating system may use the NFS protocol to communicate.)

wherein the file access protocol is CIFS. (Col. 9 lines 12-15, state that clients using the Windows operating system may communicate using CIFS)

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 7-11 and 19-23 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509), hereafter Chen in view of Baweja et al. (US 6564229), hereafter Baweja, and further in view of Lefebvre et al. (US 2002/0010665), hereafter Lefebvre.

Regarding **claim 1**, Chen (US 7076509) discloses:

In a data processing network including a client (Fig. 5, Clients 560a and/or 506b) and a file server (Fig. 5, Multi-protocol Storage Appliance 500), a method of access to a storage object (vdisk or LUN, Col 10 line 1) in the file server, said method comprising: the client using a block level access protocol (Col 9 Lines 40-42, "clients 560 generally use block-based access protocols") over the network (Fig. 5, 585 or 565) to access the storage object; and the file server accessing the storage object by accessing a file (Col 10, lines 56-60, the file server accesses the vdisk to access the data on the volume) containing data of the storage object.

which includes the file server copying the file concurrent with the client using the block level access protocol over the network to write data to the storage object (Col. 13, lines 43-49 describe the inherent snapshot capabilities of Chen's invention, and point to incorporated reference Hitz et al. (US 5,819,292) for more explanation: Col. 1 lines 26-32 of Hitz et al. describe that backups occur on servers with "active file systems" meaning that writing or reading activity can be simultaneous with any kind of copying.)

wherein the network is an IP network (Fig. 5, IP network 565), the client uses the block level protocol over a first TCP/IP connection over the network to access the storage object (iSCSI)

Regarding **claim 9 as applied to claim 8**, Chen (US 7076509) discloses:

which includes the file server also providing access to the storage object over the network by means of a file access protocol over the network, the file access protocol accessing the file containing the data of the storage object (Col 11, lines 15-20 it is disclosed that Chen is operable with both block-based and file-based protocols, i.e. "providing an integrated NAS and SAN appliance", as NAS appliances use file-based protocols.).

Regarding **claims 10-11 as applied to claims 8 and 9**, Chen (US 7076509) discloses:

wherein the client uses a UNIX or Linux operating system, and the file access protocol is NFS. (Col. 9 lines 15-18, state that a client using the UNIX (or Linux) operating system may use the NFS protocol to communicate.)

wherein the client uses a Windows operating system, and the file access protocol is CIFS. (Col. 9 lines 12-15, state that clients using the Windows operating system may communicate using CIFS)

Regarding **claim 20**, Chen (US 7076509) discloses:

In a data processing network including a client (Clients 560a and/or 506b) and a file server (Multi-protocol Storage Appliance 500), a method of access to a virtual direct access storage device (Col 12 Lines 55-60) in the file server, attributes and data of the virtual direct access storage device (Col 10, lines 56-60, the file server accesses the vdisk to access the data on the volume) being stored in at least one file in the file server (Col. 10, lines 43-47 states that a vdisk is a special file type.), said method comprising:

the client using a block level access protocol (Col 9 Lines 40-42) over the network to access the virtual direct access storage device in the file server, the file server responding to commands in accordance with the block level access protocol for access to the virtual direct access storage device by accessing the attributes and data of the virtual direct access storage device; and

the file server providing access over the network to the virtual block storage device in accordance with a file access protocol by accessing said at least one file in the file server. (Col. 9 lines 12-18)

Regarding **claim 17 as applied to claim 20**, Chen (US 7076509) discloses: which includes the file server copying the data of the virtual direct access storage device concurrent with the client using the block level access protocol over the

network to write data to the virtual direct access storage device. (Col. 13, lines 43-49 describe the inherent snapshot capabilities of Chen's invention, and point to incorporated reference Hitz et al. (US 5,819,292) for more explanation: Col. 1 lines 26-32 of Hitz et al. describe that backups occur on servers with "active file systems" meaning that writing or reading activity can be simultaneous with any kind of copying.)

Regarding **claim 21 as applied to claim 20**, Chen (US 7076509) discloses:

wherein the network is an IP network, and the block level access protocol is SCSI. (Col 9, lines 56-59, note that iSCSI, an implementation of a block level access protocol (SCSI) over TCP (TCP/IP, or an IP network)

Regarding **claims 22-23 as applied to claim 20**, Chen (US 7076509) discloses:

wherein the client uses a UNIX or Linux operating system, and the file access protocol is NFS. (Col. 9 lines 15-18, state that a client using the UNIX (or Linux) operating system may use the NFS protocol to communicate.)

wherein the client uses a Windows operating system, and the file access protocol is CIFS. (Col. 9 lines 12-15, state that clients using the Windows operating system may communicate using CIFS)

Chen (US 7076509) discloses all the limitations of claims 7-11 and 19-23 except for:

the step of the client pausing the step of writing of data to the storage object after a commit operation, and during the pause, the client performing the step of initiating the copying of the file

The general concept of pausing data writing and during the pause initiating the copying of the file is well known in the art as taught by Baweja. (Abstract, a user can pause a data copy operation (i.e. writing), and then can resume copying the file (i.e. initiating copying of the file))

It would have been obvious to one of ordinary skill in the art at the time of the invention to combine Chen with the general concept of pausing data writing in relation to a commit operation then initiating copying of data as taught by Baweja in order to free system resources to perform other operations (Baweja, Abstract, last sentence).

Chen and Baweja teach all the limitations of claims 7-11 and 19-23 except for:

the client initiates the step of the copying of the file by sending a command over a second TCP/IP connection;

the client performing the step of initiating the copying of the file by sending the command over the second TCP/IP connection.

wherein the first TCP/IP connection is concurrent with the second TCP/IP connection.

The general concept of a client requesting a replication of snapshot data over a TCP connection is well known in the art as taught by Lefebvre. (See [0086] which states that Snapshot replication can be requested by a client. It is inherent that this action would take place over a different TCP socket than that

used by iSCSI as they are different protocols which would use different TCP sockets to perform their tasks.)

It would have been obvious to one of ordinary skill in the art at the time of the invention to Chen and Baweja with the general concept of a client requesting a replication of snapshot data over a TCP connection as taught by Lefebvre in order to back the data snapshots more securely backed up off-site. Claims 28-29 is rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509 B1) as applied to claims 27 above, and further in view of Busser (US 2002/0095616 A1) and in further view of Hashemi (US 6934804).

Regarding claims 28-29,

Chen et al. (US 7076509 B1) discloses all the limitations of claim 28 except for that the attributes of the device specify an internal organization that includes a RAID level.

The general concept of storing RAID level information in metadata files is well known in the art as taught by Busser ([0029] "The disk metadata 100 contains data that the controller uses to assist in the operation and management of the RAID system.").

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen et al. (US 7076509 B1) with the general concept of storing RAID level information in metadata files in order to make information about the RAID level more available to users of the file system.

Chen, Chen, and Busser teach all the limitations of claims 28-29 except that the internal organization described includes striping.

The general concept of including striping data as part of RAID data is well known in the art as taught by Hashemi. (Col. 11, lines 3-4 "RAID attributes like those of striping or redundancy")

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen, Chen, and Busser with the general concept of including striping data as part of RAID data as taught by Hashemi in order to support all kinds of RAID configurations.

Claim 33 is rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509 B1) as applied to claim 27 above, and further in view of Chen et al. (US 7010553 B2).

Regarding claim 33,

Chen et al. (US 7076509 B1) discloses all the limitations of claim 33 except for the network file server including an IP replication facility.

The general concept of backing files up over a network is well known in the art as taught by Chen et al. (US 7010553 B2) (Col. 3 lines 14-20 "a mirror is established and stored in a remote site").

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen et al. (US 7076509 B1) with the general concept of backing files up over a network as taught by Chen et

al. (US 7010553 B2) in order to "improve reliability and facilitate disaster recovery." ('553 Col. 3 lines 14-15)

Claims 40 and 45-48 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2).

Regarding claims **40 and 45-48**, Chen et al. (US 7076509 B1) discloses:

A network file server (Fig. 6, 600) comprising:

data storage; (Disks connected to Disk driver 650)

an interface for coupling the data storage to a data network; (Media access 610) and

at least one processor programmed for permitting clients in the data network to access the data storage in accordance with a plurality of access protocols; (inherent)

the data storage containing at least one file for storing file attributes and for storing metadata defining a virtual direct access storage device and for storing data of the virtual direct access storage device; (File system 665 and/or VDISK module 670)

the access protocols including at least one block level access protocol for access to the virtual direct access storage device by accessing the metadata and data of the virtual direct access storage device (FC 630, iSCSI 628); and

the access protocols including at least one file access protocol for accessing said at least one file. (NFS 620, CIFS 622)

wherein the metadata includes attributes of the virtual direct access storage device, and the attributes of the virtual direct access storage device and the data of the virtual direct access storage device are stored together in a single file. (Col. 10, lines 43-47 states that a vdisk is a special file (singular) type.)

wherein the metadata includes attributes of the virtual direct access storage device, and the attributes of the virtual direct access storage device and the data of the virtual direct access storage device are stored in separate files in a common file system. (Col 10 lines 47-56 describe that an embodiment of a vdisk may have one file (stream) containing metadata and other files in the file system (streams) would have data.)

wherein the metadata includes attributes of the virtual direct access storage device and the attributes of the virtual direct access storage device specify an internal organization of the virtual direct access storage device. (See Fig. 7, which shows the various ways metadata, is stored in a file to describe the vdisk.)

wherein the at least one file access protocol includes NFS. (Col. 9 lines 12-18)

wherein the at least one file access protocol includes CIFS. (Col. 9 lines 12-18)

wherein the block-level access protocol includes the iSCSI protocol. (Col 9, lines 56-59)

wherein the block-level access protocol includes the SCSI protocol.

(Col 9, lines 40-42)

which includes a snapshot copy facility for creating snapshot copies of said at least one file. (Col. 13 lines 43-46 disclose a snapshot system)

wherein the attributes define an internal organization of the virtual direct access storage device and those attributes are stored in the single file. (Fig. 5 of fully incorporated application 10/216453 (now US 7107385) shows that the metadata does give an internal organization of the VLUN, showing where particular data is actually stored on a physical device.)

Chen et al. (US 7076509 B1) discloses all the limitations of claims 37-40 and 45-48 except for a facility for remote replication of at least one file over an IP network and that the remote replication facility is used to replicate snapshots.

The general concept of a facility for backing files up over a network is well known in the art as taught by Chen et al. (US 7010553 B2) (Col. 3 lines 14-20 "a mirror is established and stored in a remote site" In addition, '553 teaches that the file backed up maybe a snapshot file in Col. 3 lines 25-30).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen et al. (US 7076509 B1) with the general concept of backing files up over a network as taught by Chen et al. (US 7010553 B2) in order to "improve reliability and facilitate disaster recovery." ('553 Col. 3 lines 14-15)

Claims 41-42 is rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) as applied to claims 40 above, and further in view of Busser and in further view of Hashemi (US 6934804).

Regarding claims 41-42,

Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) disclose all the limitations of claim 28 except for that the attributes of the device specify an internal organization that includes a RAID level.

The general concept of storing RAID level information in metadata files is well known in the art as taught by Busser ([0029] “The disk metadata 100 contains data that the controller uses to assist in the operation and management of the RAID system.”).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) with the general concept of storing RAID level information in metadata files in order to make information about the RAID level more available to users of the file system.

Chen, Chen, and Busser teach all the limitations of claims 41-42 except that the internal organization described includes striping.

The general concept of including striping data as part of RAID data is well known in the art as taught by Hashemi. (Col. 11, lines 3-4 “RAID attributes like those of striping or redundancy”)

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen, Chen, and Busser with the general concept of including striping data as part of RAID data as taught by Hashemi in order to support all kinds of RAID configurations.

Claims 31-32, 43-44 and 49 are rejected under 35 U.S.C. 103(a) as being unpatentable over Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) as applied to claims 25, 30, and 40 above, and further in view of Bolosky (US 2001/0052021).

Regarding claims 31-32, 43-44 and 49,

Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) teach all the limitations of claims 43-44 and 49 except for the initiation of the replication being done by the client over a second concurrent TCP/IP connection.

The general concept of a client initiating a command over a separate concurrent TCP/IP connection is well known in the art as taught in Bolosky (Abstract, a data link for control information, and a data link for data transfer are taught)

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Chen et al. (US 7076509 B1) and Chen et al. (US 7010553 B2) with the general concept of a remote client issuing a command over a separate TCP/IP connection as taught by Bolosky in order to facilitate efficient and useful communications between clients and servers. (Bolosky, [0002] last sentence).

(10) Response to Argument

Rejections of claims 7-11 and 19-23 under 35 U.S.C. 112 2nd Paragraph

The arguments regarding these claims (Brief, pp 17-19) are moot because the Examiner has withdrawn the rejection of these claims.

Claims 27, 30, 24-26

Regarding claims 27, 30, and 24-26, Appellant argues that Chen et al. (7076509, hereafter '509) fails to disclose that the attributes of the virtual direct access storage device, which are stored in a single file with the data of the virtual direct access device, "include a specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access storage device from the single file to the data storage". (Brief, pp. 20-24)

The examiner disagrees, as cited in the above rejection, Fig. 5 of the fully incorporated application 10/216453 (now US 7107385) shows information being stored in a single file that indicates an internal organization of the virtual direct access storage device. Further, in Col. 12 line 48 - Col. 14 line 34 of US 7107385 which is a description of Fig. 5, there are multiple items that are a 'specification of an internal organization' of the virtual disk that provide information relative to the mapping of the data so that it can be retrieved from the actual storage device. For one example, in Col. 13 lines 15-30 show that the file contains information relative to mapping the data in the virtual disk to the data structure. ("metadata section 522 that includes a type field 524 and an xinode field that references another on-disk inode structure"). Appellant implies on page 22 of the brief that an internal organization is limited to 'a RAID level or a striping

scheme or related information', however, in Appellant's specification these are only listed as examples of information about an 'internal organization', it is not an explicitly limiting definition of the term. (see Appellant's specification p. 11, lines 13-18) Since claims 27, 30, and 24-26 do not specifically recite that an internal organization is limited to those examples, the claims are not limited in that way because terms in the claims are given their broadest reasonable interpretation in light of the specification; but limitations in the specification are not read into the claims.

Furthermore, even if Appellant's arguments regarding Chen/Rajan not disclosing the limitation that the vdisk file does not "include a specification of an internal organization of the virtual direct access storage device for mapping of the data of the virtual direct access storage device from the single file to the data storage" is correct, the claim is using intended use language to describe **how** the internal organization specification is to be used, not describing or limiting the type or kind of information that is stored within the single file. The information provided in Fig. 5 of Rajan can be used to map data to a data storage device, therefore meets the intended use limitation " for mapping of the data of the virtual direct access storage device from the single file to the data storage".

Claims 8-11 and 20-23

Regarding claims 8-11 and 20-23, the Examiner agrees as cited above that Chen does not explicitly disclose the client initiating the step of the first file server replicating the snapshot copy of the file over the network to the second file server by sending a command over a second TCP/IP connection to the first file server. Nor does Chen

disclose the client pausing the step of writing data to the storage object in the first file server after a commit operation, and during the pause, the client performing the step of initiating the step of the first file server replicating the snapshot copy of the file from the first server over the network to the second file server by sending the command over the second TCP/IP connection. This is stated in the rejection of these claims above. (Brief, pp. 25-28)

The Examiner disagrees with Appellant's assertion that it is not inherent for the client to request replication using a different protocol or a different TCP/IP connection. First, the management system 140 in Lefebvre is an SQL based management system, see [0072], thus, the requests would not be made using the iSCSI protocol as used in Chen, therefore the combination would -require- a separate TCP socket connection to the server than is used for the iSCSI data communication as the two protocols operate over different system port numbers. This is further evidenced by two definitionary references included with the Appeal Brief. The first is the IANA list of registered ports, which shows that iSCSI uses port 860 (IANA pg. 32) and that SQL managers use port 1186 (IANA pg. 41). The second evidentiary reference (a definition of a socket taken from PC Magazine) shows that a TCP/IP socket can be bound to only one port, thus connections to two different services with different ports -require- multiple socket connections.

Appellant argues on page 32 of the brief that one of ordinary skill in the art at the time of the invention would not be motivated by Baweja and Lefebvre to modify Chen. However, the Examiner has provided motivations for one of ordinary skill in the art at

the time of the invention to do so. Further, the Examiner notes that this is a combination of elements that are known in the art as shown by the Examiner, and their combination does not produce any unexpected results.

Claims 7 and 19

Regarding claims 7 and 19, the Examiner disagrees with Appellant's statement that the fact that the connections are concurrent distinguishes from Lefebvre. (Brief, pp. 33-34) As cited in the rejection of these claims, Chen, via fully incorporated reference Hitz (US 5,819,292) in Col. 1 lines 26-32 discloses backups occur on servers with "active file systems", meaning that writing and reading can be simultaneous with copying. Therefore, the connection suggested by Lefebvre to request a backup may happen concurrently with other users accessing the file system.

Claims 28-29

Regarding claims 28 and 29, Appellant argues that Chen/Rajan teach away from including information about RAID and striping levels in the metadata. The Examiner disagrees, Busser teaches storing information about a RAID level of an array as cited in the above rejection, and while Rajan may be directed to allowing a user to request the creation of a vdisk and not require any further user involvement, this does not mean that there -can not- be any further user involvement, and one of ordinary skill in the art at the time of the invention would be motivated to make as much information about the underlying data storage known as could be possible so as to allow management programs access to the information. Regarding Hashemi, specifically in Col. 11 lines 36-47 it is disclosed that virtual disks (i.e. LUNs) are assigned certain stripes/slices, and it

would be imperative that each virtual disk would have recorded upon it what slices it was to be accessing so as to prevent it from accessing a disk area not assigned to it.

Claim 33

Appellant has not provided any other substantive arguments other than those refuted for claim 27 above, and, for the sake of brevity, the Board is respectfully referred to the above. The Examiner also notes that motivation has been provided for the combination of Chen '553 and Chen '509 in the rejection of claim 33 above.

Claims 40 and 45-48

Appellant has not provided any other substantive arguments other than those refuted for claims 27 and 33 above, and, for the sake of brevity, the Board is respectfully referred to the above.

Claims 41 and 42

Appellant has not provided any other substantive arguments other than those refuted for claims 28 and 29 above, and, for the sake of brevity, the Board is respectfully referred to the above.

Claims 31-32, 43-44, and 49

The Arguments above with respect to the Chen references above apply equally here to the arguments made by Appellants to the Chen references. Further, Appellant argues that because the environment of Bolosky's server is different than the environment of Appellant's server. Further, the Appellant appears to argue that inherency was used in arguing that Bolosky teaches using two concurrent connections, one to issue commands to a server, and a separate connection to send data. However,

these limitations are taught by the abstract of Bolosky "These multiple network connections may include a control link connection for passing control information and a data funnel connection for passing data of multiple media."

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

MEK 7/29/2008

Conferees:

/Joseph E. Avellino/

Primary Examiner, Art Unit 2146

/Nathan J. Flynn/

Supervisory Patent Examiner, Art Unit 2143